

Simplification and integration in computing and cognition: the SP theory and the multiple alignment concept

J Gerard Wolff*

30 November 2012

Abstract

The main purpose of this article is to describe potential benefits and applications of the SP theory, a unique attempt to simplify and integrate ideas across artificial intelligence, mainstream computing and human cognition, with information compression as a unifying theme. The theory, including a concept of multiple alignment, combines conceptual simplicity with descriptive and explanatory power in several areas including representation of knowledge, natural language processing, pattern recognition, several kinds of reasoning, the storage and retrieval of information, planning and problem solving, unsupervised learning, information compression, and human perception and cognition.

In the *SP machine*—an expression of the SP theory which is currently realised in the form of computer models—there is potential for an overall simplification of computing systems, including software.

As a theory with a broad base of support, the SP theory promises useful insights in many areas and the integration of structures and functions, both within a given area and amongst different areas. There are potential benefits in natural language processing (with potential for the understanding and translation of natural languages), the need for a versatile intelligence in autonomous robots, computer vision, intelligent databases, maintaining multiple versions of documents or web pages, software engineering, criminal investigations, the management of big data and gaining benefits from it, the semantic web, medical

*Dr Gerry Wolff, BA (Cantab), PhD (Wales), CEng, MBCS (CITP); CognitionResearch.org, 18 Penlon, Menai Bridge, Anglesey, LL59 5LR, UK; jgw@cognitionresearch.org; +44 (0) 1248 712962; +44 (0) 7746 290775; www.cognitionresearch.org.

diagnosis, the detection of computer viruses, the economical transmission of data, and data fusion. Further development of these ideas would be facilitated by the creation of a high-parallel, web-based, open-source version of the SP machine, with a good user interface. This would provide a means for researchers to explore what can be done with the system and to refine it.

Keywords: information compression, artificial intelligence, perception, cognition, representation of knowledge, learning, pattern recognition, natural language processing, reasoning, planning, problem solving.

1 Introduction

Since about 1987, I have been fascinated by the idea that it might be possible to create a theory that simplifies and integrates ideas across artificial intelligence, mainstream computing and human cognition, with information compression as a unifying theme. The first rather crude attempt at such a theory, with the main focus on computing, was described in Wolff (1990). In the following years, progressively more refined versions of the ‘SP’ theory have been described in several peer-reviewed journal articles and, in some detail, in a book (Wolff, 2006a).¹

The main purpose of this article is to describe potential benefits and applications of the theory. In brief:

- The SP theory combines conceptual simplicity with descriptive and explanatory power.
- It has potential to simplify computing systems, including software.
- As a theory with a broad base of support, the SP theory promises to provide useful insights in many areas and is likely to facilitate the simplification and integration of structures and functions, both within a given area and amongst different areas.

These points will be discussed, each in its own section, below. But first, as a background and context to the discussion, the next section considers a little philosophy and outlines the SP theory, the multiple alignment concept, and associated ideas.

¹Bibliographic details of relevant publications may be found via links from www.cognitionresearch.org/sp.htm.

2 Background and context

2.1 A little philosophy

Amongst alternative criteria for the success of a theory, what appears to be the most widely accepted is the principle (‘Occam’s Razor’) that a good theory should combine simplicity with descriptive or explanatory power.² This equates with the idea—which echoes the underlying theme of the SP theory itself³—that a good theory should compress empirical data via a reduction of ‘redundancy’ in the data (thus increasing its ‘simplicity’), whilst retaining as much as possible of its non-redundant descriptive ‘power’. As John Barrow has written: “Science is, at root, just the search for compression in the world.” (Barrow, 1992, p. 247).

These principles are prominent in most areas of science: the Copernican heliocentric theory of the Sun and the planets is rightly seen to provide a simpler and more powerful account of the facts than Ptolemy’s Earth-centred theory, with its complex epicycles; physicists have a keen interest in the possibility that quantum theory and relativity might be unified; biology would be greatly impoverished without modern understandings of evolution, genetics, and DNA; and so on.

But in research in computer science, including artificial intelligence, there is a fragmentation of the field into a myriad of concepts and many specialisms, with little in the way of an overarching theory to pull everything together.⁴ And some areas, such as computer vision, appear to be overburdened with over-complex mathematics.⁵

Perhaps the Turing machine model of computing (Turing, 1936) provides what is needed? That theory has of course been brilliantly successful, providing theoretical underpinnings for the many useful things that can be done with computers. But it does not solve the problem of fragmentation and, although Alan Turing recognised that computers might become intelligent

²According to this criterion, a good theory in a given area lies somewhere between things which are simple but very general—such as the laws of arithmetic—and over-specific ‘theories’ that simply repeat what it is they are meant to explain.

³Hence the name ‘SP’.

⁴Similar things were said by Allen Newell about research in human perception and cognition, in his well-known essay on why “You can’t play 20 questions with nature and win” (Newell (1973). See also Newell (1992, 1990)).

⁵Of course, mathematics is often very useful—the SP theory has a solid mathematical foundation—but large amounts of mathematics suggest that the ideas are on the wrong track, somewhat like Ptolemy’s epicycles. And the limitations of mathematical notations can be an obstacle to the development of other concepts that do not fit easily into a mathematical perspective.

(Turing, 1950), the Turing theory, in itself, does not tell us how!

Whether or not the SP theory succeeds in plugging these gaps, I believe there is a pressing need, in computer science and AI, for a stronger focus on the simplification and integration of ideas.⁶ In all areas, including quite specialised areas, a theory that scores well in terms of simplicity and power is, compared with any weaker theory, likely to yield deeper insights and better integration of concepts, both within a given area and amongst different areas (see also Section 5).

2.2 Outline of the SP theory

2.2.1 Origins and motivation

Much of the inspiration for the SP theory is a body of research, pioneered by Fred Attneave (see, for example, Attneave, 1954), Horace Barlow (see, for example, Barlow, 1959, 1969), and others, showing that many aspects of the workings of brains and nervous systems may be understood as compression of information. For example, when we view a scene with two eyes, the image on the retina of the left eye is almost exactly the same as the image on the retina of right eye, but our brains merge the two images into a single percept, and thus compress the information (Barlow, 1969).⁷

The theory also draws on principles of ‘minimum length encoding’ pioneered by Solomonoff (1964); Wallace and Boulton (1968); Rissanen (1978), and others. The close connection that exists between information compression and concepts of probability (see, for example, Li and Vitányi, 2009) is an important motivation for the theory.

These ideas, with several observations about computing and AI, suggested that it might be possible to develop a theory bridging artificial intelligence, mainstream computing and human cognition, with information compression as a unifying theme (Wolff, 2006a, Chapter 2).

2.2.2 Elements of the theory

The main elements of the SP theory are:

- All knowledge is expressed as arrays of atomic symbols (*patterns*) in one or two dimensions.

⁶Notice that collaboration amongst different specialisms—which is necessary or useful for some kinds of projects—does not in itself achieve simplification and integration of ideas.

⁷This focus on compression of information in binocular vision is distinct from the more usual interest in the way that slight differences between the two images enables us to see the scene in depth.

- Patterns that represent data coming in from the environment are classified as ‘New’. Patterns representing stored knowledge are classified as ‘Old’.
- The system is designed for the unsupervised learning of Old patterns by compression of New patterns.
- An important part of this process is the economical encoding of New patterns in terms of Old patterns, if they exist. This may be seen to achieve such things as pattern recognition, parsing or understanding of natural language, or other kinds of interpretation of incoming information in terms of stored knowledge, including several kinds of reasoning.
- Compression of information is achieved via the matching and unification (merging) of patterns, with an improved version of dynamic programming (Wolff, 2006a, Appendix A) providing flexibility in matching, and with key roles for the frequency of occurrence of patterns, and their sizes.⁸
- The concept of *multiple alignment*, outlined in Section 2.2.5, is a powerful central idea, similar to the concept of multiple alignment in bioinformatics but with important differences.
- Owing to the intimate connection, already mentioned, between information compression and concepts of probability, the SP theory provides for the calculation of probabilities for inferences made by the system, and probabilities for parsings, recognition of patterns, and so on.

2.2.3 The SP62 and SP70 computer models

The SP theory is currently expressed in two computer models, SP62 and SP70. Although SP62 is a subset of SP70 which can build multiple alignments but which lacks any capability for unsupervised learning, it has been convenient to maintain them as separate models.

The SP models and their precursors have played a key part in the development of the theory:

- By forcing precision in the definition of the theory where there might otherwise be vagueness.

⁸In the SP theory, the matching and unification of patterns is seen as being closer to the bedrock of information compression, and with greater heuristic value, than more mathematical techniques such as wavelets, arithmetic coding, or the like. The mathematics which provides the basis for those techniques may itself be founded on the matching and unification of patterns (see Wolff, 2006a, Chapter 10).

- By providing a convenient means of encoding the simple but important mathematics that underpins the SP theory, and performing relevant calculations, including calculations of probability.
- By providing a means of seeing quickly the strengths and weaknesses of proposed mechanisms or processes.
- By providing a means of demonstrating what can be achieved with the theory.

2.2.4 The SP machine

The SP models may be regarded as implementations of the *SP machine*, an expression of the SP theory and a means for it to be applied.

A useful step forward in the development of the SP theory and the SP machine would be the creation, as an open-source software virtual machine, of a high-parallel, web-based version, taking advantage of powerful search mechanisms in one of the existing search engines, and with a good user interface. This would provide a means for researchers to explore what can be done with the system and to refine it.

Further ahead, there may be a case for the creation of new hardware, dedicated to the building of multiple alignments and other processes in the SP framework.

2.2.5 The multiple alignment concept

A concept of multiple alignment is widely used in bioinformatics where it means a process of arranging, in rows or columns, two or more DNA sequences or amino-acid sequences so that matching symbols—as many as possible—are aligned orthogonally in columns or rows.

In the SP theory, this concept has been borrowed and adapted, as illustrated in Figures 1 and 2. The main difference between the two concepts is that, in bioinformatics, all sequences have the same status, whereas in the SP theory, the system attempts to create a multiple alignment which enables one New pattern⁹ (shown in row 0 in Figure 1 and column 0 in Figure 2) to be encoded economically in terms of one or more Old patterns (shown in rows 1 to 8 in Figure 1 and in columns 1 to 4 in Figure 2).

In Figure 1, the New pattern is a sentence, the Old patterns represent grammatical structures (including words), and the whole multiple alignment may be seen as a parsing of the sentence in terms of a grammar.

⁹In some applications, there may be more than one New pattern in a multiple alignment.

0				t w o				k i t t e n s					p l a y	0
1								< Nr 5 k i t t e n >						1
2								< N Np < Nr						2
3				< D Dp 4 t w o >										3
4				< NP < D				> < N						4
5													< Vr 1 p l a y >	5
6														6
7	< S Num			; < NP									> < V	> > 7
8	Num PL ;							Np					Vp	8

Figure 1: The best multiple alignment created by the SP62 model with a store of Old patterns like those in rows 1 to 8 (representing grammatical structures, including words) and the New pattern (representing a sentence to be parsed) shown in row 0.

Notice how, in row 8, the system marks the dependency between the plural subject of the sentence (‘Np’) and the plural main verb (‘Vp’).

Although this is not illustrated here, the system also supports the production of language as well as its analysis, it facilitates the integration of syntax and semantics, and it is robust in the face of errors in the sentence to be parsed or the stored grammatical patterns, or both.

As an indication of the versatility of the multiple alignment concept, Figure 2 illustrates how it may model recognition of an unknown creature at multiple levels of abstraction, as described in the caption to the figure.

Not illustrated here is the way the framework can accommodate cross-classification and the integration of class hierarchies with part-whole hierarchies, and how recognition may be achieved despite errors in incoming data or stored knowledge, or both.

These are just two of several things that can be modelled within the multiple alignment framework. Others are summarised in Section 3, below.

2.2.6 A possible neural basis for the SP theory

A tentative part of the SP theory is the idea that the cortex of the brains of mammals—which is, topologically, a two-dimensional sheet—may be, in some respects, like a sheet of paper on which neural analogues of SP patterns may be written—as shown schematically in Figure 3. Unlike information written on a sheet of paper, there are neural connections between patterns—as shown in the figure—and communications amongst them.

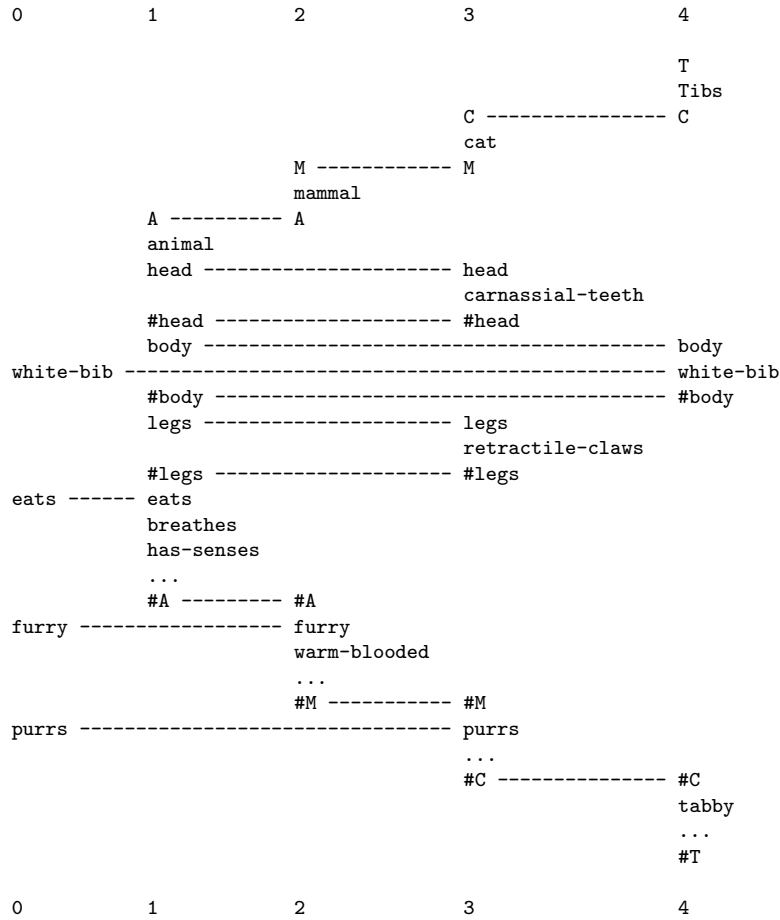


Figure 2: A multiple alignment showing how an unknown creature, with features shown in the New pattern in column 0, may be recognised in terms of Old patterns at multiple levels of abstraction: as an animal (column 1), a mammal (column 2), a cat (column 3), and as a specific individual, ‘Tibs’ (column 4).

2.2.7 Related research

To my knowledge, the SP theory is unique in attempting to simplify and integrate ideas across artificial intelligence, mainstream computing, and human cognition.

But there has for some time been an interest in unified theories of cognition such as SOAR (see, for example, Laird, 2012) and ACT-R (see, for example, Anderson et al., 2004).¹⁰ And, although it is still very much a minority interest amongst researchers in artificial intelligence, there has in

¹⁰Both of them inspired by Allen Newell (eg, Newell, 1973).

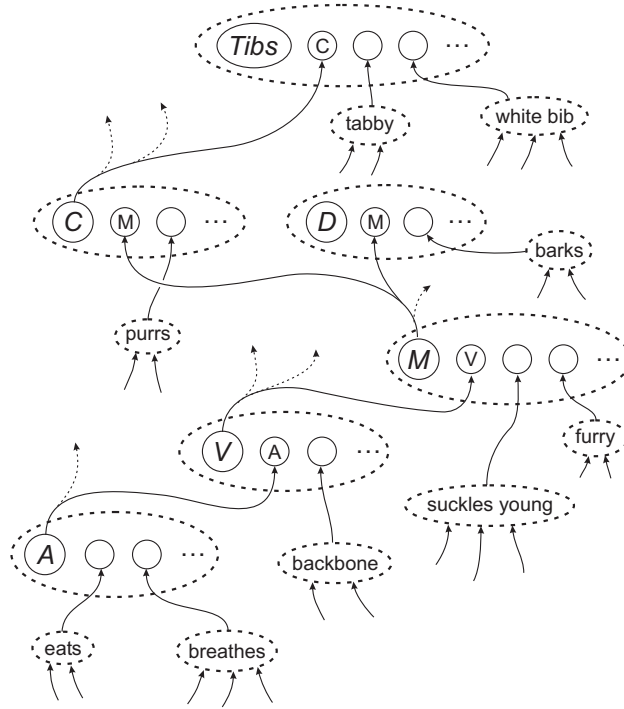


Figure 3: Schematic representation of hypothesised neural analogues of SP patterns and their inter-connections. Reproduced from Figure 11.6 of Wolff (2006a), with permission.

recent years been renewed interest in artificial intelligence as a whole (see, for example, Schmidhuber et al., 2011).

As mentioned in Section 2.2.1, the SP theory is inspired in part by research showing the importance of information compression in brains and nervous systems, and it draws on other research on principles of minimum length encoding, and the close relation that exists between information compression and concepts of probability.

2.2.8 Unfinished business

Like most theories, the SP theory cannot claim to solve every problem in its target areas. The main shortcomings of the theory as it stands now are:

- *Processing of information in two or more dimensions.* No attempt has yet been made to generalise the SP models to work with patterns in two dimensions, although that appears to be feasible to do, as outlined in Wolff (2006a, Section 13.2.1). As noted in Wolff (2006a, Section 13.2.2), it is possible that information with dimensions higher than two may

be encoded in terms of patterns in one or two dimensions, somewhat in the manner of architects' drawings. With the aid of triangulation to determine depths, a 3D structure may be stitched together from several partially-overlapping 2D views, in much the same way that, in digital photography, a panoramic view may be created from partially-overlapping pictures (For reviews of related work in computer vision, see Prince, 2012; Szeliski, 2011).

- *Recognition of perceptual features in speech and visual images.* For the SP system to be effective in the processing of speech or visual images, it seems likely that some kind of preliminary processing will be required to identify low level perceptual features such as, in the case of speech, phonemes, formant ratios, or formant transitions, or, in the case of visual images, edges, angles, colours, luminances, or textures. In vision, at least, it seems likely that the SP framework itself will prove relevant since edges may be seen as zones of non-redundant information between uniform areas containing more redundancy and, likewise, angles may be seen to provide significant information where straight edges, with more redundancy, come together (See, for example, Attneave, 1954). As a stop-gap solution, the preliminary processing may be done using existing techniques for the identification of low-level perceptual features (see, for example, Prince, 2012, Chapter 13).
- *Unsupervised learning.* A limitation of the SP70 model as it is now is that it cannot learn intermediate levels of abstraction in grammars or the like, and it cannot learn discontinuous dependencies in natural language syntax or other kinds of information. I believe these problems are soluble and that solving them will greatly enhance the capabilities of the system for the unsupervised learning of structure in data.
- *Processing of numbers.* The SP models work with atomic symbols such as ASCII characters or strings of characters with no intrinsic meaning. In itself, the SP system does not recognise the arithmetic meaning of numbers such as '37' or '652' and will not process them correctly. However, the system has the potential to handle mathematical concepts correctly, if it is supplied with the relevant mathematical rules (Wolff, 2006a, Chapter 10). As a stop-gap solution, existing technologies may provide whatever arithmetic processing may be required.

3 The SP theory combines simplicity with descriptive and explanatory power

In principle, it should be possible to evaluate scientific theories quite precisely in terms of simplicity and power, and to compare one theory with another, by measuring their ability to compress empirical data. But techniques and technologies are not yet well enough developed to make this feasible, and there is in any case the difficulty that it is rare for rival theories to address precisely the same body of empirical data. So in evaluating theories we still need to rely on more-or-less informal judgements of simplicity and explanatory or descriptive power.

In those terms, the SP theory appears to score well:

- The SP theory, including the multiple alignment concept, is not trivially simple but it is not unduly complicated either. The SP70 model, which is the most comprehensive expression of the theory as it stands now, is embodied in an ‘exec’ file requiring less than 500 KB of storage space.
- Largely because of the versatility of the multiple alignment concept, the SP theory can model a good range of concepts and phenomena in artificial intelligence, mainstream computing, and human cognition, as summarised in the following subsections.

3.1 The SP theory as a theory of computing

In Wolff (2006a, Chapter 4), I have argued that the operation of a Post canonical system (Post, 1943) may be interpreted in terms of the SP theory and that, since it is accepted that the concept of a Post canonical system is equivalent to the Turing machine model of computing (which is itself widely regarded as a definition of the concept of ‘computing’), the SP theory may itself be regarded as a theory of computing—and not merely a theory of human cognition or artificial intelligence. By contrast with the Turing model and the Post canonical system, the fundamentals of the SP theory are multiple alignment and the compression of information by the matching and unification of patterns.

It is pertinent to mention here that the SP theory also has what I believe to be some useful things to say about the nature of mathematics (see Wolff, 2006a, Chapter 10).

3.2 Representation of knowledge

Despite the simplicity of representing knowledge with patterns, the way they are processed within the multiple alignment framework gives them the versatility to represent several kinds of knowledge, including grammars for natural languages, class hierarchies, part-whole hierarchies, discrimination networks, relational tuples, if-then rules, and others mentioned elsewhere in this article. This versatility is achieved without recourse to the kinds of constructs used in, for example, the ‘OWL 2’ web ontology language:¹¹ ‘SubClassOf’, ‘ObjectProperty’, ‘SubObjectPropertyOf’, and so on.

One universal format for knowledge and one universal framework for processing means that any subset of the different kinds of knowledge may be combined flexibly and seamlessly according to need.

3.3 Natural language processing

One of the main strengths of the SP system is in natural language processing:

- As illustrated in Figure 1, grammatical rules, including words and their grammatical markers, may be represented with SP patterns.
- Both the parsing and production of natural language may be modelled via the building of multiple alignments.
- The framework provides an elegant means of representing discontinuous dependencies in syntax, including overlapping dependencies such as number dependencies and gender dependencies in languages like French.¹²
- The system may also model non-syntactic ‘semantic’ structures such as class-inclusion hierarchies and part-whole hierarchies (as noted in Section 3.2 and elsewhere).
- The system supports the seamless integration of syntax with semantics, providing a role for non-syntactic knowledge in the processing of natural language, and giving potential for the understanding of natural languages and interlingua-based translations amongst languages.

¹¹See, for example, the *OWL 2 Web Ontology Language Primer* at www.w3.org/TR/owl2-primer.

¹²As an illustration, row 8 of Figure 1 enforces the rule that a plural subject for a sentence (‘Np’) must be followed by a plural main verb (‘Vp’).

- The system is robust in the face of errors of omission, commission or substitution in sentences to be analysed, or stored linguistic knowledge, or both.
- The importance of context in the processing of language (see, for example, Iwanska and Zadrozny, 1997) is accommodated in the way the system searches for a global best match for patterns: any pattern or partial pattern may be a context for any other.

3.4 Pattern recognition

The system provides a powerful framework for pattern recognition:

- It can model pattern recognition at multiple levels of abstraction, as outlined in Section 2.2.5.
- As noted in Section 3.2, the system also provides for cross-classification and the integration of class-inclusion relations with part-whole hierarchies.
- The system provides for seamless integration of pattern recognition with various kinds of reasoning (Section 3.5). As a simple example, if Tibs is recognised as a mammal on the strength of his or her furry coat, we may infer from the relevant pattern that the creature is warm blooded.
- A probability may be calculated for any given classification or any associated inference.
- The system is robust in the face of errors of omission, commission or substitution in incoming data, or stored knowledge, or both.
- As in the processing of natural language (Section 3.3), the importance of context in recognition (see, for example, Oliva and Torralba, 2007) is accommodated in the way the system searches for a global best match for patterns.

3.5 Reasoning

The SP system can model several kinds of reasoning including one-step ‘deductive’ reasoning, abductive reasoning, reasoning with probabilistic decision networks and decision trees, reasoning with ‘rules’, nonmonotonic reasoning and reasoning with default values, reasoning in Bayesian networks (including

‘explaining away’), causal diagnosis, and reasoning which is not supported by evidence.

Since these several kinds of reasoning all flow from one computational framework (multiple alignment), they may be seen as aspects of one process, working individually or together without awkward boundaries.

Plausible lines of reasoning may be achieved, even when relevant information is incomplete.

Probabilities of inferences may be calculated, including extreme values (0 or 1) in the case of logic-like ‘deductions’.

3.6 Information storage and retrieval

The SP theory provides a versatile model for database systems, with the ability to accommodate object-oriented structures (with class hierarchies and inheritance of attributes, including multiple inheritance), as well as relational ‘tuples’, and network and tree models of data (Wolff, 2007). Unlike most database systems, it provides a simple and elegant means of combining class hierarchies with part-whole hierarchies.

It lends itself most directly to information retrieval in the manner of query-by-example but it appears to have potential to support the use of query languages such as SQL.

3.7 Planning and problem solving

The SP framework provides a means of planning a route between two places, and, with the translation of geometric patterns into textual form, it can solve the kind of geometric analogy problem that may be seen in some puzzle books and IQ tests (Wolff, 2006a, Chapter 8).

3.8 Unsupervised learning

The SP70 model can derive a plausible grammar from a set of sentences without supervision or error correction by a ‘teacher’, without the provision of ‘negative’ samples, and without the grading of samples from simple to complex. It thus overcomes restrictions on what can be achieved with some other models of learning (eg, Gold, 1967), and reflects more accurately what is known about how children learn their first language or languages (Wolff, 1988).

The model draws on earlier research showing that inductive learning via principles of ‘minimum length encoding’ (the matching and unification of

patterns) can lead to the discovery of entities—such as words in natural languages or ‘objects’ in the world of vision—that are psychologically natural, as illustrated in Figure 4. This principle—the discovery of natural structures via information compression (DONSVIC)—is also likely to apply to the discovery of classes of entity.

As was mentioned in Section 2.2.8, the model, as it stands now, is not able to derive intermediate levels of abstraction or discontinuous dependencies in data, but those problems appear to be soluble.



Figure 4: Part of a parsing created by program MK10 (Wolff, 1977) from a 10,000 letter sample of English (book 8A of the Ladybird Reading Series) with all spaces and punctuation removed. The program derived this parsing from the sample alone, without any prior dictionary or other knowledge of the structure of English. Reproduced from Figure 7.3 of Wolff (1988), with permission.

3.9 Compression of information

Since information compression is central to the workings of the SP system, it is natural to consider whether the system might provide useful insights in that area. The potential of the system to detect and encode discontinuous dependencies in data suggests that it may be able to extract kinds of redundancy in information that are not accessible via standard methods for the compression of information.

In terms of the trade-off that exists between computational resources that are required and the level of compression that can be achieved, it seems likely that the system will be towards the ‘up market’ end of the spectrum—by contrast with LZW algorithms and the like, which have been designed to be ‘quick-and-dirty’, sacrificing performance for speed on low-powered computers.

3.10 Human perception and cognition

Since much of the inspiration for the SP theory has come from evidence, mentioned in Section 2.2, that, to a large extent, the workings of brains and nervous systems may be understood in terms of information compression, the theory is about perception and cognition as well as artificial intelligence and mainstream computing.

That said, the main elements of the theory—the multiple alignment concept in particular—are theoretical constructs derived from what appears to be necessary to model, in an economical way, such things as pattern recognition, reasoning, and so on. In Wolff (2006a, Chapter 11), I have described in outline how such things as SP patterns and multiple alignments may be realised with neurons and connections between them. Given the enormous complexity of the brains of humans and other mammals, and given the difficulties of working with living brains, it is likely to be difficult to obtain direct evidence for or against the kinds of neural structures that I have described. But, because of the strength of support for the constructs in the SP theory, I believe that area of research deserves attention.

4 How the SP theory may promote simplification of computing systems

Apart from the potential of the SP theory to simplify concepts in artificial intelligence and mainstream computing (Section 3), it can help to simplify computing systems, including software.

The principle to be described here is already familiar in the way databases and expert systems are structured.

Early databases were each written as a monolithic system containing everything needed for its operation, and likewise for expert systems. But it soon became apparent that, in the case of databases, a lot of effort could be saved by creating a generalised ‘database management system’ (DBMS), with a user interface and mechanisms for the storage and retrieval of information, and creating new databases by loading the DBMS with different bodies of information, according to need. In a similar way, an ‘expert system shell’, with a user interface and with mechanisms for inference and for the storage and retrieval of information, eliminates the need to create those facilities repeatedly in each new expert system.

The SP system takes this principle further. It aims to provide a general-purpose ‘intelligence’—chiefly the multiple alignment framework—and thus save the need to create those kinds of mechanisms repeatedly in different AI applications: pattern recognition, natural language processing, several kinds of reasoning, planning, problem solving, unsupervised learning, and more.

The principle extends to conventional applications, since the matching and unification of patterns, and a process of searching amongst alternative matches for patterns (which are elements of the multiple alignment framework), are used in many kinds of application outside the world of artificial intelligence.

The way in which this principle may help to simplify computing systems is shown schematically in Figure 5.

In a conventional computer, shown at the top of the figure, there is a central processing unit (‘CPU’) which is relatively simple and without ‘intelligence’.

Software in the system, shown to the right, is a combination of two things:

- Domain-specific knowledge such as knowledge of accountancy, geography, the organisation and procedures of a business, and so on.¹³
- Processing instructions to provide the intelligence that is missing in the CPU, chiefly processes for the matching and unification of patterns, for searching amongst alternative matches to find one or more that are ‘good’, and the implicit formation of multiple alignments. These kinds of processing instructions—let us call them ‘MUP’ instructions—are required in many different kinds of application, meaning that there

¹³Even in systems such as DBMSs or expert-system shells, mentioned above, the domain-specific knowledge which is loaded into the system may be regarded as part of the software.

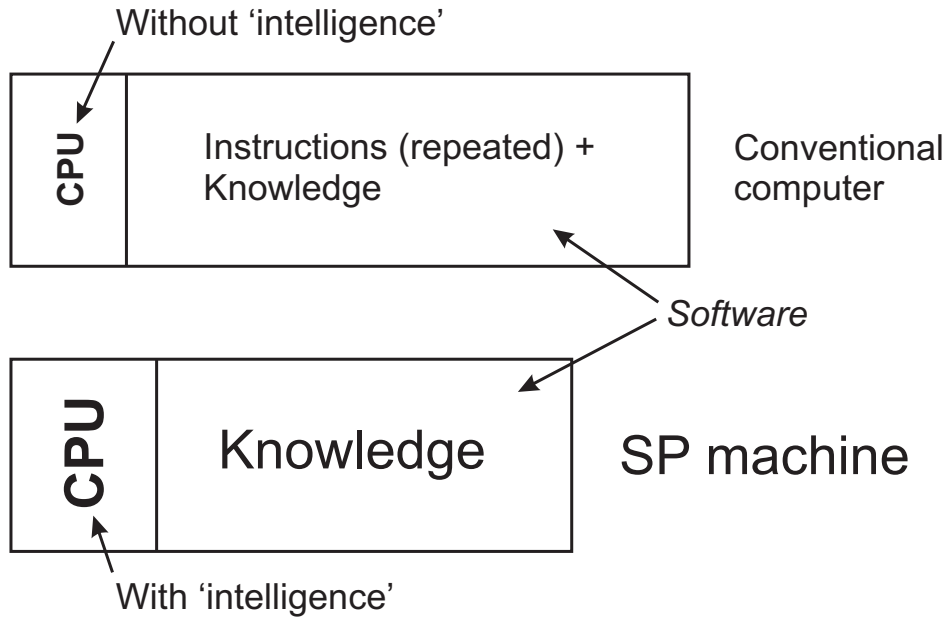


Figure 5: Schematic representations of a conventional computer and an SP machine, as discussed in the text.

is a considerable amount of redundancy in conventional computing systems—where the term ‘system’ includes both the variety of programs that may be run, as well as the hardware on which they run.

In the SP machine, shown schematically at the bottom of the figure, the CPU is more complicated (as suggested by the larger size of ‘CPU’ in the figure) and provides the ‘intelligence’ that is missing in the CPU of a conventional computer. This should mean that MUP instructions may be largely eliminated from the software, leaving just the domain-specific knowledge. The increased complexity of the CPU should be more than offset by cutting out redundancy in the software, meaning an overall reduction in complexity—as suggested by the relatively small size in the figure of the SP machine, compared with the conventional computer.

5 Potential benefits of the SP theory in some areas of application

This section describes some areas of application for the SP theory, and its potential benefits, drawing on the ‘enabling technologies’ summarised in Section 3. These examples are probably just the tip of the iceberg. To mix our

metaphors, there is a rich vein to be mined.

Because of its importance, I believe it is worth repeating the principle, mentioned in Section 2.1, that, even in specialised areas, a theory like the SP theory that has a broad base of support is likely to prove useful.

There seem to be two main reasons for this:

- A theory that scores well in terms of simplicity and power is likely to provide deeper insights than others. The germ theory of disease serves us better than the idea that diseases are caused by “bad air”; recognising the role of oxygen in combustion is more practical than the phlogiston theory; and so on.
- Even in one relatively narrow area of application, a broadly-supported theory is likely to promote the simplification and integration of structures and functions, both within the given area and between that area and others.

Even if existing technologies in a given area are doing well, there is likely to be a case for examining how the SP theory may be applied, and its potential benefits.

5.1 Applications in the processing of natural language

As noted in Section 3.3, the SP system supports the parsing and production of language, it provides an elegant means of representing discontinuous dependencies in syntax, it is robust in the face of errors in data, it facilitates the integration of syntax with semantics, and it provides a role for context in the processing of language. There are many potential applications, including such things as checking for grammatical errors in text; predictive typing that takes account of grammar as well as the spellings of individual words; natural-language front-ends to databases; and the understanding and translation of natural languages.

It is true that some of these things are already done quite well with existing technologies but, as suggested earlier, there may nevertheless be advantages in exploring what can be done with SP system. It is likely, for example, to facilitate the integration of syntax with semantics, and it is likely to smooth the path for the integration of natural language processing with other aspects of intelligence: reasoning, learning, pattern recognition, and so on.

5.1.1 Crowdsourcing as a route towards the understanding and translation of natural languages?

Given the apparent simplicity of representing syntactic and non-syntactic knowledge with patterns (Section 3.2), it may be feasible to use crowdsourcing in the manner of Wikipedia to create syntactic rules for at least one natural language and a realistically large set of ontologies to provide associated meanings.

With this kind of knowledge on the web and with a high-parallel, on-line version of the SP machine (as sketched in Section 2.2.4), it may be feasible to achieve understanding of natural language, at least at an elementary level.

With the same ontologies and with syntactic knowledge for two or more languages, it may be feasible to achieve translations between languages, with the ontologies functioning as an ‘interlingua’.¹⁴

When translations are required amongst several different languages, this approach has the attraction that it is only necessary to create n mappings, one between each of the n languages and the interlingua. This contrasts with syntax-to-syntax systems such as Google Translate, where $n!/2(n-2)!$ pairs of languages are needed.

Another possible advantage is that, arguably, machine translation using an interlingua has the potential to produce translations of a better quality than may be achieved with syntax-to-syntax translation.

5.2 Towards a versatile intelligence for autonomous robots

If a robot is to survive and be effective in a challenging environment like Mars, where communication lags restrict the help that can be provided by people, there are advantages if it can be provided with human-like intelligence and versatility—as much as possible. A theory like the SP theory, which provides a unified view of several different aspects of intelligence, is a good candidate for the development of that versatile intelligence.¹⁵

¹⁴See, for example, ‘Interlingual machine translation’, Wikipedia, bit.ly/RlTcHU.

¹⁵This reference to autonomous robots does not in any way endorse or defend the irresponsible use of airborne ‘drones’ or other autonomous robots for military or related kinds of operations, including the use of any such device without robust safeguards and controls.

5.3 Computer vision

Despite the problems noted in Section 2.2.8, I believe the SP system has potential in the area of computer vision:

- It has potential to simplify and integrate several areas in computer vision, including feature detection and alignment, segmentation, deriving structure from motion, stitching of images together, stereo correspondence, scene analysis, and object recognition (see, for example, Szeliski, 2011). With regard to the last two topics:
 - Since scene analysis is, in some respects, similar to the parsing of natural language, and since the SP system performs well in parsing, it has potential in scene analysis as well. In the same way that a sentence may be parsed successfully without the need for explicit markers of the boundaries between successive words or between successive phrases, a scene may be analysed into its component parts without the need for explicit boundaries between objects or other elements in a scene.
 - The system provides a powerful framework for pattern recognition, as outlined in Section 3.4. It seems likely that this can be generalised for object recognition.
- The SP system has potential for unsupervised learning of the knowledge required for recognition. For example, discrete objects may be identified by the matching and merging of patterns within stereo images (*cf.* Marr and Poggio (1979))¹⁶ or within successive frames in a video, in much the same way that the word structure of natural language may be discovered via the matching and unification of patterns (Section 3.8). The system may also learn such things as classes of entity, and associations between entities, such as the association between black clouds and rain.
- The system is likely to facilitate the seamless integration of vision with other aspects of intelligence: reasoning, planning, problem solving, natural language processing, and so on.
- As noted earlier, the system is robust in the face of errors of omission, commission or substitution—an essential feature of any system that is to achieve human-like capabilities in vision.

¹⁶As with information compression, a focus on the isolation of discrete objects in binocular vision is distinct from the more usual interest in the way that slight differences between the two images enables us to see the scene in depth.

- There are potential solutions to the problems outlined in Section 2.2.8:
 - It is likely that the framework can be generalised to accommodate patterns in two dimensions.
 - As noted earlier, 3D structures may be modelled using 2D patterns, somewhat in the manner of architects' drawings. Knowledge of such structures may be built via the matching and unification of partially-overlapping 2D views.
 - The framework has potential to support both the discovery and recognition of low-level perceptual features.

5.4 A versatile model for intelligent databases

As described in Wolff (2007), the SP system provides a versatile framework for the creation of intelligent databases:

- It has the versatility to support the storage and retrieval of several kinds of knowledge, including object-oriented structures with inheritance of attributes (including multiple inheritance), relational tuples, and data that is structured as networks or trees.
- As noted earlier, it provides for the integration of class-inclusion hierarchies with part-whole hierarchies.
- It is free from the anomalies and inefficiencies that arise in existing systems that use a relational database as a means of creating object-oriented or entity-relationship data structures (see Wolff, 2007, Section 4.2.3).
- It provides for retrieval of information in the manner of query-by-example but has the potential to support the creation of query languages like SQL.
- Unlike conventional databases, it supports various kinds of reasoning, and other kinds of 'intelligence'.

5.5 Maintaining multiple versions of a document or web page

The multiple alignment concept provides an elegant means of maintaining multiple versions of any document or web page, such as, for example, versions

in different languages. Parts which are shared amongst different versions—such as pictures or diagrams—may be kept separate from parts that vary amongst different versions—such as text in a specific language. As with other kinds of data, the framework provides the means of maintaining hierarchies of classes, with cross classification if that is required, and of maintaining part-whole hierarchies and their integration with class hierarchies.

Updates to high-level classes appear directly in all lower-level classes without the need for information to be copied amongst different versions.

5.6 Software engineering: object-oriented design, verification and validation

For the following main reasons, the SP system has potential in software engineering:

- Databases provide the mainstay of many software engineering projects and, as we have seen (Section 5.4), the SP system has potential in that area.
- As noted in Section 5.4, the SP system provides for object-oriented design with inheritance of attributes, including multiple inheritance, and it provides ‘intelligence’ as well.
- It appears that the SP system may serve to model real-world procedures in business, administration, the law, and so on—although this has not yet been examined in any detail.

As was suggested in Section 4, the SP system has potential to reduce or eliminate ‘MUP’ instructions and their repetition in different applications. To the extent that that proves possible, we may reduce the opportunities for errors to be introduced. There is potential to reduce or eliminate the need for ‘verification’ in software development, with corresponding improvements in the quality of software.

The SP machine also has potential in ‘validation’: helping to ensure that what is being developed is what is required. If domain-specific knowledge, including the requirements for a software system, can be loaded directly into an SP machine, without the need for the traditional kind of programming, this will help to ensure that what the system does is what the users want. There is potential for users to have tighter control than is traditional over what a system does, with more freedom to make changes when required.¹⁷

¹⁷Notice that the distinction, just referenced, between domain-specific knowledge and MUP instructions is different from the often-made distinction between declarative

5.7 Management of data and the drawing of inferences in criminal investigations

It seems likely that an industrial-strength version of the SP machine would be useful in crime investigations:

- It would provide a means of storing and managing the data that are gathered in such investigations, often in large amounts.
- It has potential to recognise similarities between a given crime and other crimes, either current or past—and likewise for suspects.
- It may prove useful in piecing together coherent patterns from partially-overlapping fragments of information, in much the same way that partially-overlapping digital photographs may be stitched together to create a larger picture.
- Given the potential of the SP system for different kinds of reasoning, including reasoning with information that is not complete, the system may prove useful in suggesting avenues to be explored, perhaps highlighting aspects of an investigation that might otherwise be overlooked.

5.8 Managing ‘big data’ and gaining value from it

Because information compression is central to the workings of the SP system, and because the system has potential to deliver relatively high levels of compression (Section 3.9), it may prove useful in reducing the sizes of the very large data sets known as ‘big data’, and thus helping to make them more manageable.

The capabilities of the SP system in pattern recognition are potentially useful in the process of searching big data for particular kinds of information.

With solutions to the residual problems in unsupervised learning (Section 2.2.8, it seems likely that the SP machine, via the ‘DONSVIC’ principle (Section 3.8), will prove useful in discovering ‘interesting’ or ‘useful’ structures in big data, including significant entities and classes of entity, and rules or regularities, including discontinuous dependencies in data.

Since the concept of multiple alignment in the SP theory is a modified version of that concept in bioinformatics, the SP machine may prove useful in the discovery or detection of interesting structures in DNA sequences or amino-acid sequences.

programming and imperative or procedural programming. It is envisaged that the SP machine would reduce or eliminate the MUP instructions but that it would be able to model both the declarative ‘what’ and the procedural ‘how’ of the real world.

5.9 Knowledge, reasoning, and the semantic web

The SP framework may make a useful contribution to the development of the ‘semantic web’—a ‘web of data’ to provide machine-understandable meaning for web pages (see, for example, Berners-Lee et al., 2001; Shadbolt et al., 2006).¹⁸ In this connection, the main attractions of the system appear to be:

- *Simplicity, versatility and integration in the representation of knowledge.* As noted in Section 3.2, the SP system combines simplicity in the underlying format for knowledge with the versatility to represent several different kinds of knowledge—and it facilitates the seamless integration of different kinds of knowledge.
- *Versatility in reasoning.* The system provides for several kinds of reasoning and their inter-working (Section 3.5).
- *Potential for natural language understanding.* As noted in Section 3.3, seamless integration of syntactic and semantic knowledge in the SP system gives it potential for the understanding of natural language. If that proves possible, semantic structures may be derived from textual information in web pages, without the need, in those web pages, for separate provision of ontologies or the like.
- *Potential for automatic learning of ontologies.* As was suggested in Section 5.8, the SP system has potential, via the ‘DONSVIC’ principle, for the extraction of interesting or useful structures from big data. These may include the kinds of ontologies which have been a focus of interest in work on the development of the semantic web.
- *Uncertainty and vagueness.* The SP system is inherently probabilistic and it is robust in the face of incomplete information and errors of commission or substitution. These capabilities appear promising as a means of coping with uncertainty and vagueness in the semantic web (Lukasiewicz and Straccia, 2008).

5.10 Medical diagnosis

The way in which the SP system may be applied in medical diagnosis is described in Wolff (2006b). The potential benefits of the SP system in that area of application include:

- A format for representing diseases that is simple and intuitive.

¹⁸See also www.w3.org/standards/semanticweb/.

- An ability to cope with errors and uncertainties in diagnostic information.
- The simplicity of storing statistical information as frequencies of occurrence of diseases.
- The system provides a method for evaluating alternative diagnostic hypotheses that yields true probabilities.
- It is a framework that should facilitate the unsupervised learning of medical knowledge and the integration of medical diagnosis with other AI applications.

The main emphasis in Wolff (2006b) is on medical diagnosis as pattern recognition. But the SP system has potential for causal reasoning (see Wolff, 2006a, Section 7.9) which, in a medical context, may enable the system to make probabilistic inferences such as “The patient’s fatigue may be caused by anaemia which may be caused by a shortage of iron in the diet”.

5.11 Detection of computer viruses

The detection of already-known computer viruses and other malware can be more subtle and difficult than simply looking for exact matches for the ‘signatures’ of viruses. The offending code may be rather similar to perfectly legitimate code and it may be contained within a compressed (‘packed’) executable file.

Here, the SP system and its capabilities for pattern recognition and learning may prove useful:

- Recognition in the SP system is probabilistic, it does not depend on the presence or absence of any particular feature or combination of features, and it can cope with errors.
- By compressing known viruses into a set of SP patterns, the system can reduce the amount of information needed to specify viruses and, as a consequence, it can reduce the amount of processing needed for the detection of viruses.

5.12 Support for the economical transmission of data

A compressed version of a body of information, I , may be seen to comprise two parts:

- A ‘grammar’ for I containing patterns, at one or more levels of abstraction, that occur repeatedly in I .
- An ‘encoding’ of I in terms of the grammar, including non-repeating information in I .

Where there is significant redundancy in I —which is true of most kinds of natural-language text and most kinds of images—the encoding is likely to be much smaller than the grammar. A grammar for I may also support the economical encoding of other bodies of information, provided they contain the same kinds of structures as I .

These things can provide the means of transmitting information very economically:

- A receiver (such as a TV set) may be equipped with a grammar for the kind of information it is designed to receive (eg, TV images), and some version of the SP system so that it may decode incoming information in terms of the grammar.
- Instead of transmitting ‘raw’ data, or even data that has been compressed in the usual manner (containing both the grammar and the encoding), the encoding by itself would be sufficient.
- The original data may be reconstructed fully, without any loss of information, by decoding the transmitted information in terms of the stored grammar.

A simple version of this idea, using a dictionary of words as a ‘grammar’ for ordinary text, is already recognised (see, for example, Giltner et al. (1983), Storer (1988, Chapter 1)). The potential advantage of the SP system is the creation of more sophisticated grammars yielding higher levels of compression, and the application of the idea to kinds of information—such as images—where an ordinary dictionary would not work.

Since ordinary TV works well without this kind of mechanism, some people may argue that there is no need for anything different. But:

- The growing popularity of video, TV and films on mobile services is putting pressure on mobile bandwidth.¹⁹

¹⁹See, for example, “Data jam threat to UK mobile networks”, BBC News, 2012-11-16, bbc.in/T5g14s.

- It is likely that some of the bandwidth for terrestrial TV will be transferred to mobile services,²⁰ creating an incentive to use the remaining bandwidth efficiently.
- There would be benefits in, for example, the transmission of information from a robot on Mars, or any other situation where a relatively large amount of information needs to be transmitted as quickly as possible over a relatively low-bandwidth channel.

5.13 Data fusion

There are many situations where it can be useful to merge or integrate two or more sources of information, normally when there is some commonality amongst the different sources. For example, in ecological research there may be a need to integrate different sources of information about the movement of whales or other creatures around the world; and in the management of a website, there may be a need to consolidate a set of web pages that are saying similar things.

At first sight, the SP system appears to be tailor made for this area of application. After all, the merging of fully or partially matching patterns lies at the heart of the SP system. And for certain kinds of data fusion projects, that kind of capability could indeed prove to be effective and useful.

However, for better results, it is likely that solutions will be needed for the problems outlined in Section 2.2.8. And, to rival human capabilities in the consolidation of natural language texts, the potential of the system for the understanding of natural language would need to be developed to the point where it could understand that two bodies of text are saying the same thing but with different words.

6 Conclusion

The potential benefits of the SP theory and SP machine include simplification and integration of concepts in computer science (including artificial intelligence), and simplification of computing systems, including software.

As a theory with a broad base of support, the SP theory is likely to provide useful insights in many areas and to promote the integration of structures and functions, both within a given area and amongst different areas.

²⁰See, for example, “TVs will need retuning again to make room for mobile services”, The Guardian, 2012-11-16, bit.ly/ZZ9N8C.

The SP theory promises new and improved solutions to problems in several areas including natural language processing (with potential for the understanding and translation of natural languages), the need for a versatile intelligence in autonomous robots, computer vision, intelligent databases, maintaining multiple versions of documents or web pages, software engineering, criminal investigations, the management of big data and gaining benefits from it, the semantic web, medical diagnosis, the detection of computer viruses, the economical transmission of data, and data fusion.

A useful step forward in the development of these ideas would be the creation, as a software virtual machine, of a high-parallel, web-based, open-source version of the SP machine, taking advantage of powerful search mechanisms in one of the existing search engines, and with a good user interface. This would provide a means for researchers to explore what can be done with the system and to refine it.

References

- J. R. Anderson, D. Bothell, M. D. Byrne, S. Douglass, C. Lebiere, and Y. Qin. An integrated theory of the mind. *Psychological Review*, 111(4):1036–1060, 2004.
- F. Attneave. Some informational aspects of visual perception. *Psychological Review*, 61:183–193, 1954.
- H. B. Barlow. Sensory mechanisms, the reduction of redundancy, and intelligence. In HMSO, editor, *The Mechanisation of Thought Processes*, pages 535–559. Her Majesty’s Stationery Office, London, 1959.
- H. B. Barlow. Trigger features, adaptation and economy of impulses. In K. N. Leibovic, editor, *Information Processes in the Nervous System*, pages 209–230. Springer, New York, 1969.
- J. D. Barrow. *Pi in the Sky*. Penguin Books, Harmondsworth, 1992.
- T. Berners-Lee, J. Hendler, and O. Lassila. The Semantic Web. *Scientific American*, pages 35–43, May 2001.
- M.E. Giltner, J. C. Mueller, and R. R. Fiest. Data compression, encryption, and in-line transmission system. United States Patent number 4386416, 1983. Filed: 2 June 1980; Issued: 31 May 1983.
- M. Gold. Language identification in the limit. *Information and Control*, 10: 447–474, 1967.

- L. Iwanska and W. Zadrozny. Introduction to special issue on context in natural language processing. *Computational Intelligence*, 13(3):301–308, 1997.
- J. E. Laird. *The Soar Cognitive Architecture*. The MIT Press, Cambridge, Mass., 2012. ISBN-13: 978-0-262-12296-2.
- M. Li and P. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications*. Springer, New York, 2009.
- T. Lukasiewicz and U. Straccia. Managing uncertainty and vagueness in description logics for the Semantic Web. *Proceedings of the Royal Society, Series B*, 6(4):291–308, 2008.
- D. Marr and T. Poggio. A computational theory of human stereo vision. *Proceedings of the Royal Society of London. Series B*, 204(1156):301–328, 1979.
- A. Newell. You can’t play 20 questions with nature and win: projective comments on the papers in this symposium. In W. G. Chase, editor, *Visual Information Processing*, pages 283–308. Academic Press, New York, 1973.
- A. Newell, editor. *Unified Theories of Cognition*. Harvard University Press, Cambridge, Mass., 1990.
- A. Newell. Précis of *Unified Theories of Cognition*. *Behavioural and Brain Sciences*, 15(3):425–437, 1992.
- A. Oliva and A. Torralba. The role of context in object recognition. *Trends in cognitive sciences*, 11(12):520–527, 2007.
- E. L. Post. Formal reductions of the general combinatorial decision problem. *American Journal of Mathematics*, 65:197–268, 1943.
- S. J. D. Prince. *Computer Vision: Models, Learning, and Inference*. Cambridge University Press, Cambridge, 2012. ISBN 978-1-107-01179-3.
- J. Rissanen. Modelling by the shortest data description. *Automatica-J, IFAC*, 14:465–471, 1978.
- J. Schmidhuber, K. R. Thórisson, and M. Looks, editors. *Proceedings of the Fourth International Conference on Artificial General Intelligence (AGI 2011)*, volume 6830 of *Lecture Notes in Artificial Intelligence*, 2011. Springer. ISBN 978-3-642-22886-5.

- N. Shadbolt, W. Hall, and T. Berners-Lee. The semantic web revisited. *IEEE Intelligent Systems*, 21(3):96–101, 2006.
- R. J. Solomonoff. A formal theory of inductive inference. Parts I and II. *Information and Control*, 7:1–22 and 224–254, 1964.
- J. A. Storer. *Data Compression: Methods and Theory*. Computer Science Press, Rockville, Maryland, 1988.
- R. Szeliski. *Computer Vision: Algorithms and Applications*. Springer, London, 2011. ISBN 978-1-84882-934-3.
- A. M. Turing. On computable numbers with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society*, 42:230–265 and 544–546, 1936.
- A. M. Turing. Computing machinery and intelligence. *Mind*, 59:433–460, 1950.
- C. S. Wallace and D. M. Boulton. An information measure for classification. *Computer Journal*, 11(2):185–195, 1968.
- J. G. Wolff. The discovery of segments in natural language. *British Journal of Psychology*, 68:97–106, 1977. See: www.cognitionresearch.org/lang_learn.html#wolff_1977.
- J. G. Wolff. Learning syntax and meanings through optimization and distributional analysis. In Y. Levy, I. M. Schlesinger, and M. D. S. Braine, editors, *Categories and Processes in Language Acquisition*, pages 179–215. Lawrence Erlbaum, Hillsdale, NJ, 1988. See: www.cognitionresearch.org/lang_learn.html#wolff_1988.
- J. G. Wolff. Simplicity and power—some unifying ideas in computing. *Computer Journal*, 33(6):518–534, 1990. See: www.cognitionresearch.org/papers/early/early.htm.
- J. G. Wolff. *Unifying Computing and Cognition: the SP Theory and Its Applications*. CognitionResearch.org, Menai Bridge, 2006a. ISBNs: 0-9550726-0-3 (ebook edition), 0-9550726-1-1 (print edition). Distributors, including Amazon.com, are detailed on www.cognitionresearch.org/books/sp_book/retailers.htm. The publisher and its website was previously CognitionResearch.org.uk.

- J. G. Wolff. Medical diagnosis as pattern recognition in a framework of information compression by multiple alignment, unification and search. *Decision Support Systems*, 42:608–625, 2006b. See: www.cognitionresearch.org/papers/applications/medical/medical_applications.htm.
- J. G. Wolff. Towards an intelligent database system founded on the SP theory of computing and cognition. *Data & Knowledge Engineering*, 60:596–624, 2007. See: www.cognitionresearch.org/papers/dbir/dbir.htm.